# Trainable & Dynamic Computing: Error Backpropagation Through Physical Media

Michiel Hermans, Michael Burm, Joni Dambre, Peter Bienstman

July 25, 2014

arXiv:1407.6637v1 [cs.NE] 24 Jul 2014

**Machine learning algorithms, and more in particular neural networks, arguably experience a revolution in terms of performance. Currently, the best systems we have for speech recognition, computer vision and similar problems are based on neural networks, trained using the half-century old backpropagation algorithm. Despite the fact that neural networks are a form of analog computers, they are still implemented digitally for reasons of convenience and availability. In this paper we demonstrate how we can design physical linear dynamic systems with non-linear feedback as a generic platform for dynamic, neuro-inspired analog computing. We show that a crucial advantage of this setup is that the error backpropagation can be performed physically as well, which greatly speeds up the optimisation process. As we show in this paper, using one experimentally validated and one conceptual example, such systems may be the key to providing a relatively straightforward mechanism for constructing highly scalable, fully dynamic analog computers.**

In a variety of forms, neural networks have seen an exponential rise in attention the last decade. Once deemed an unworkable academic curiosity, neural networks trained with gradient descent are now outperforming other, more classical approaches in a broad number of challenging benchmarks. One dramatic example is their state-of-the-art performance in computer vision [15]. This, including many other examples, are problems that are considered easy for humans, and hard for conventional computer algorithms. On top of a number of semi-heuristic methods, this result has been obtained using the *backpropagation* algorithm, a method that has been around since the 60's.

One highly interesting set of neural architectures are so-called recurrent neural networks (RNN), neural networks that have temporal feedback loops. This makes them highly suited for problems that have a natural time element, such as speech recognition [9] and character-based language modelling [21, 11]. Conventional neural networks may be limited in solving such tasks, as they can only include a finite temporal context (usually a window of data of a fixed duration). RNNs, on the other hand, can – at least in principle – have indefinitely long memory. Indeed, it is believed that temporal feedback loops are primary functional components of the human brain. Unfortunately, RNNs are also slow to train compared to more common feedforward neural networks, as they cannot benefit from parallel computing architectures as much. This means that RNNs have only seen a limited impact of the boost of research into neural networks of the last decade.

Currently, (recurrent) neural networks are implemented on digital devices, mostly for reasons of convenience and availability. At their core, however, they are analog computers, and they come far closer to mimicking the workings of the brain than classic computation algorithms do. A computer made up of presently available components that would have enough processing power to match that of the human brain is estimated to consume about 100 Megawatts of power (according to estimates of IBM research), a full seven orders of magnitude more than the 20 Watts required by our own brains. If we ever wish to achieve any degree of scalability for devices performing brain-like computation, we will need to embrace physical realisations of analog computers, where information is encoded by real-valued physical variables, and where these data are processed by letting them interact through physical nonlinear dynamic systems.

One line of research that has been partially successful in accomplishing this goal is that of Reservoir Computing (RC) [17]. This paradigm, which combines several previous lines of research [13, 18], essentially employs randomly constructed dynamic systems to process time series. The dynamic system is fed with a time-varying signal that needs to be processed. If the dynamic system is sufficiently varied, non-linear and high-dimensional, it will act as an efficient random feature generator, which expands the input signal into a high-dimensional space in which the time series processing problem becomes much easier to solve. All that needs to be optimised is a linear mapping of these features to a desired output signal, which can be performed efficiently with any linear algebra solver.

The last decade of research into RC has shown a variety of interesting examples of physical implementations of analog processors. It has been shown to work with water ripples [5], mechanical constructs [3, 10], electro-optical devices [16, 20], fully optical devices [2] and nanophotonic circuits [24, 25]. Despite some successes, the RC concept still faces the problem of being ineffective for tasks that require a great deal of modelling power (e.g., natural language processing, video stream processing, etc.). The main reason being that, due to the random initialisation of the employed systems, any particular feature that needs to be extracted from the input data needs to be present in the system

by sheer luck. Even for moderately demanding tasks this requires a disproportionally large dimensionality of the system at hand, up to tens of thousands of state variables in practice [22]. When the input dimensionality grows, the odds of having the correct features present becomes vanishingly small. By contrast, neural networks trained by the backpropagation algorithm can build the required nonlinear features internally during the training process.

In recent work we have shown that it is possible to extend backpropagation to models of physical dynamic systems [12]. We showed that it can serve as an efficient automated tool to find highly non-trivial solutions for complex dynamic problems. So far, this work relies on simulation, however. The backpropagation algorithm operates on a computer model of the dynamic system. If we would use it to train physical analog computers, we would face the same scaling issues that we encountered with systems that approach the human brain. In this paper, we offer a definition of a dynamic system that encompasses a very broad set of analog computing systems, including all forms of conventional neural networks. Crucially, we show that the backpropagation algorithm can be performed physically on such systems with only minor additional requirements, which strongly reduces the external computational demands of the optimisation process, hence greatly speeding up the training process.

# Results

We start by introducing an abbreviated notation for a multivariate convolution. If $\mathbf{x}(t)$ is a multivariate signal, and $\mathbf{W}(t)$ a matrix with time-varying elements defined for $t > 0$, we define the signal $\mathbf{y}(t)$ as the convolution of $\mathbf{x}(t)$ with $\mathbf{W}(t)$ as follows:

$$\mathbf{y}(t) = \int_0^\infty dt' \; \mathbf{W}(t')\mathbf{x}(t-t') = [\mathbf{W} * \mathbf{x}](t).$$

We will consider systems as follows. We will assume that there are a set of $N$ signal input sources $s_i(t)$ which excite the LDS, and a set of $M$ output receivers which receive an output signal $o_i(t)$. We can write both sets as a single source and receiver vector $\mathbf{s}(t)$ and $\mathbf{o}(t)$, respectively. The LDS will cause the following transformation between the source and receiver:

$$\mathbf{o}(t) = [\mathbf{W}_{\mathbf{so}} * \mathbf{s}](t), \tag{1}$$

where the impulse response, or first order Volterra kernel $\mathbf{W}_{\mathbf{so}}(t)$ characterises the transfer function of the system. Furthermore, we are able to use the LDS reversely, where the receivers act as sources, and the sources as receivers. In this case, $\mathbf{o}(t)$ acts as the input of the system and $\mathbf{s}(t)$ represent the signal receives at the places of the original sources. The LDS is described by:

$$\mathbf{s}(t) = \left[\mathbf{W}_{\mathbf{so}}^{\mathsf{T}} * \mathbf{o}\right](t), \tag{2}$$

i.e., the operation is reciprocal. One important example of such a system which we will use for the rest of the paper is the propagation of waves through a linear medium. Suppose for instance we have a chamber in which we place a set of different speakers and a set of microphones, the signals the microphones receive would indeed be described by equation 1, where $\mathbf{W}_{\mathbf{so}}(t)$ would be determined by the shape of the room, the absorption of the walls, the air density, etc. If we then replace each speaker by a microphone and vice versa, the signal we received would be described by equation 2. Another set of examples are systems described by the linear heat (or diffusion) equation. The source would have a controllable temperature, and the receivers would be thermometers, where the medium causes an operation as described by the above equations.

Linear systems can only perform linear operations on their input signal. In order for the full system to be able to model non-linear relationships, we add nonlinear feedback. We provide a set of $M_{\mathbf{a}}$ additional receivers and $N_{\mathbf{a}}$ sources. The signal that is detected at these new receivers gets sent through a non-linear operator $\mathbf{f} : \mathbb{R}^{M_{\mathbf{a}}} \to \mathbb{R}^{N_{\mathbf{a}}}$ and is fed back into the system via the new sources. We denote the signal after the function as $\mathbf{a}(t)$. The impulse response matrix for the transition from the input sources to the receivers for the nonlinear feedback we denote as $\mathbf{W}_{\mathbf{sa}}(t)$, and those for the transition from the non-linear feedback sources to the output receivers and non-linear feedback receivers with $\mathbf{W}_{\mathbf{ao}}(t)$ and $\mathbf{W}_{\mathbf{aa}}(t)$, respectively. A schematic diagram of the full system is shown in the top of Figure 1. The system is described as follows:

$$\begin{aligned} \mathbf{a}(t) &= \mathbf{f}\left([\mathbf{W}_{\mathbf{sa}} * \mathbf{s}](t) + [\mathbf{W}_{\mathbf{aa}} * \mathbf{a}](t)\right) \\ \mathbf{o}(t) &= [\mathbf{W}_{\mathbf{so}} * \mathbf{s}](t) + [\mathbf{W}_{\mathbf{ao}} * \mathbf{a}](t) \end{aligned} \tag{3}$$

As we argue in the supplementary material, these equations can easily be identified with neural networks, including all kinds of deep networks, recurrent neural networks, etc. If we wish to use this system as a trainable model for signal processing, we need to be able to define gradients for the parameters we can change. First of all we define a cost functional $C(\mathbf{o}(t))$, a function of the whole history of $\mathbf{a}(t)$ in the interval $t \in \{0 \cdots T\}$ that we wish to minimize. Next, we define the partial derivative of $C(\mathbf{o}(t))$ w.r.t. $\mathbf{o}(t)$ as $\mathbf{e}_{\mathbf{o}}(t)$. The error backpropagation process is then described by the following equations:

$$\begin{aligned} \mathbf{e}_{\mathbf{a}}(s) &= \mathbf{J}^{\mathsf{T}}(s)\left(\left[\mathbf{W}_{\mathbf{ao}}^{\mathsf{T}} * \mathbf{e}_{\mathbf{o}}\right](s) + \left[\mathbf{W}_{\mathbf{aa}}^{\mathsf{T}} * \mathbf{e}_{\mathbf{a}}\right](s)\right) \\ \mathbf{e}_{\mathbf{s}}(s) &= \left[\mathbf{W}_{\mathbf{so}}^{\mathsf{T}} * \mathbf{e}_{\mathbf{o}}\right](s) + \left[\mathbf{W}_{\mathbf{sa}}^{\mathsf{T}} * \mathbf{e}_{\mathbf{a}}\right](s). \end{aligned} \tag{4}$$

Here, $s = T - t$, i.e., the equations run backward in time (which in practice means that we play $\mathbf{e}_{\mathbf{o}}(t)$ backwards as input for the system). From the variables $\mathbf{e}_{\mathbf{a}}(t)$ and $\mathbf{e}_{\mathbf{s}}(t)$, gradients w.r.t. all impulse response matrices within the system, and w.r.t. $\mathbf{s}(t)$ can be found.

The above equations describe a continuous-time system. Most time-related machine learning problems work on discrete-time sequences, however. In order to use the system to process discrete time series, we use an input encoding scheme which was first introduced in [1]. Suppose we have an input time series $\mathbf{x}_i$ that we wish to map on an output time series $\mathbf{y}_i$. First of all we define an encoding which transforms the vector $\mathbf{x}_i$, the $i$-th instance of the input data sequence, into a continuous time signal segment $\mathbf{s}_i(t)$:

$$\mathbf{s}_i(t) = \mathbf{s}_b(t) + \mathbf{M}(t)\mathbf{x}_i \quad \text{for} \quad t \in [0 \cdots P],$$

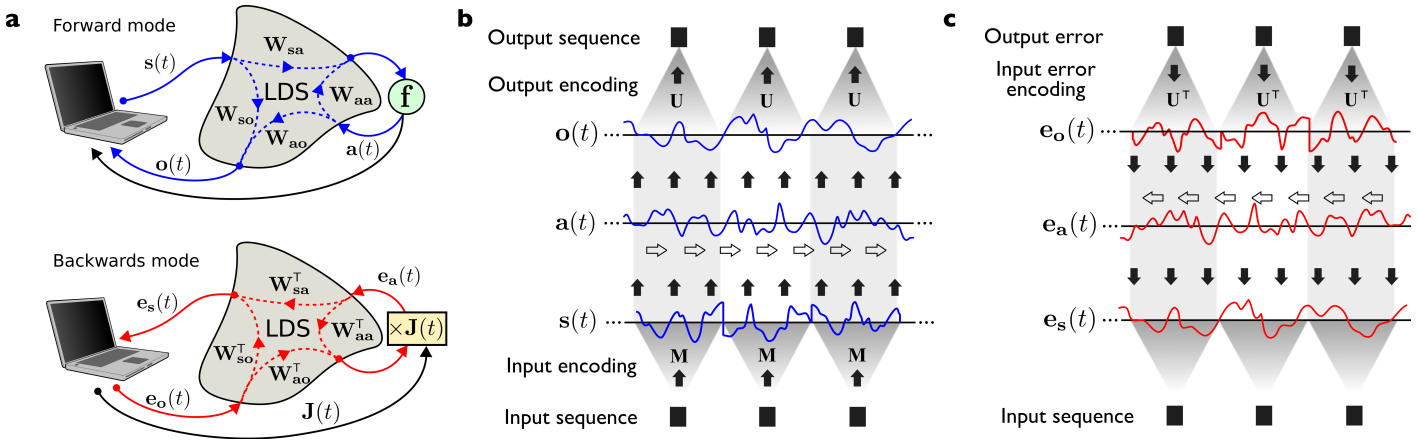


**Figure 1:** **A:** Illustration of the most general setup of the physical neural network studied in this paper. The top diagram shows how the signals propagate through the system and the nonlinear feedback with blue arrows during the forward pass. The filter operations are depicted with dashed lines running through the linear dynamic system (LDS), which is depicted as a grey blob. The bottom diagram shows the error backpropagation phase, where the signal runs backwards through all functional dependencies. Here, filter operations run in the opposite direction such that they are represented by the transpose of their impulse response matrices. Note that the computer is not in the loop during the forward or backward pass, but only serves to send out a predefined signal, and to record at the same time. **B:** Depiction of the masking principle in the forward direction. At the bottom we see three consecutive instances of an input time series. Each of these is converted into a finite time segment through the masking signals $\mathbf{M}(t)$. These segments are next concatenated in time and serve as the input signal $\mathbf{s}(t)$ for the dynamic system (where time runs according to the white arrows), and which in turn generates an output signal $\mathbf{o}(t)$. The output signal $\mathbf{o}(t)$ is divided into finite length pieces, which are decoded into output instances of an output time series using the output masks $\mathbf{U}(t)$. **C:** The backpropagation process happens in a completely similar manner as in the forward direction. This time, the transpose of the output masks serve as the encoding masks. Finally, the input error signal $\mathbf{e_s}(t)$ is also segmented in time before it is used to determine the gradients w.r.t. $\mathbf{M}(t)$

where $P$ is the masking period and $\mathbf{s}_b(t)$ is a bias time trace. The matrix $\mathbf{M}(t)$ are the so-called input masks, defined for a finite time interval of duration $P$. The input signal $\mathbf{s}(t)$ is now simply the time-concatenation of the finite time segments $\mathbf{s}_i(t)$

The output encoding works in a very similar fashion. If there is a time series $\mathbf{y}_i$ which represents the output associated with the $i$-th instance of the input time series, we can define an output mask $\mathbf{U}(t)$. We divide the time trace of the system output $\mathbf{o}(t)$ into segments $\mathbf{o}_i(t)$ of duration $P$. The $i$-th network output instance is then defined as

$$\mathbf{y}_i = \mathbf{y}_b + \int_0^P dt\ \mathbf{U}(t)\mathbf{o}_i(t),$$

with $\mathbf{y}_b$ a bias vector. The process described here is essentially a form of time multiplexing. The backpropagation phase happens very similarly. Suppose we have a time series with instances $\mathbf{e}_i$, which are the gradient of a cost function w.r.t. the output instances $\mathbf{y}_i$. Completely equivalent to the input masking we can now define the error signal $\mathbf{e_o}(t)$ as a time concatenation of finite time segments $\mathbf{e}^i_\mathbf{o}(t)$:

$$\mathbf{e}^i_\mathbf{o}(t) = \mathbf{U}^\mathsf{T}(t)\mathbf{e}_i.$$

Using this signal as input to the system during backpropagation will provide us with an error $\mathbf{e_s}(t)$, which in turn provides can be used to determine the gradient for the masking signals $\mathbf{M}(t)$. Note that the signal encoding and decoding happens on an external PC as well, and not physically. These operations can be fully parallelised, however, such that we do not lose the advantages of the processing power of the system. Note that, in practice, we always measure signals using discrete sampling. Similarly, input signals are generated with finite temporal precision. In reality this means that the input and output masking signals also will be defined as a finite set of numbers.

The training process works as follows. First we sample a certain amount of data, run it through the physical system and record the output signals. Next, we compute the cost of the output and construct the signal $\mathbf{e_o}(t)$, which we then also run through the system physically, and record the signal $\mathbf{e_s}(t)$ (and $\mathbf{e_a}(t) if applicable$). Once this is completed, we can compute gradients w.r.t. the relevant parameters (impulse response matrices if applicable, and input and output masks). We then subtract these gradients from the parameters, multiplied with a (small) learning rate. This process is then repeated for many iterations until we reach satisfactory performance or until the experiment stops.

## A real-life acoustic setup

We have tested the principles described above by building a system that uses the propagation of acoustic waves as an LDS. To reduce the complexity of the setup we work with only one signal source (a small computer speaker) and one receiver (a voice microphone). The sound enters a 6-meter long plastic tube via a paper funnel, and the microphone receives the signal at the other end of the tube. The tube will delay the signal and will introduce reflections and resonance frequencies. The received signal is electronically trun-

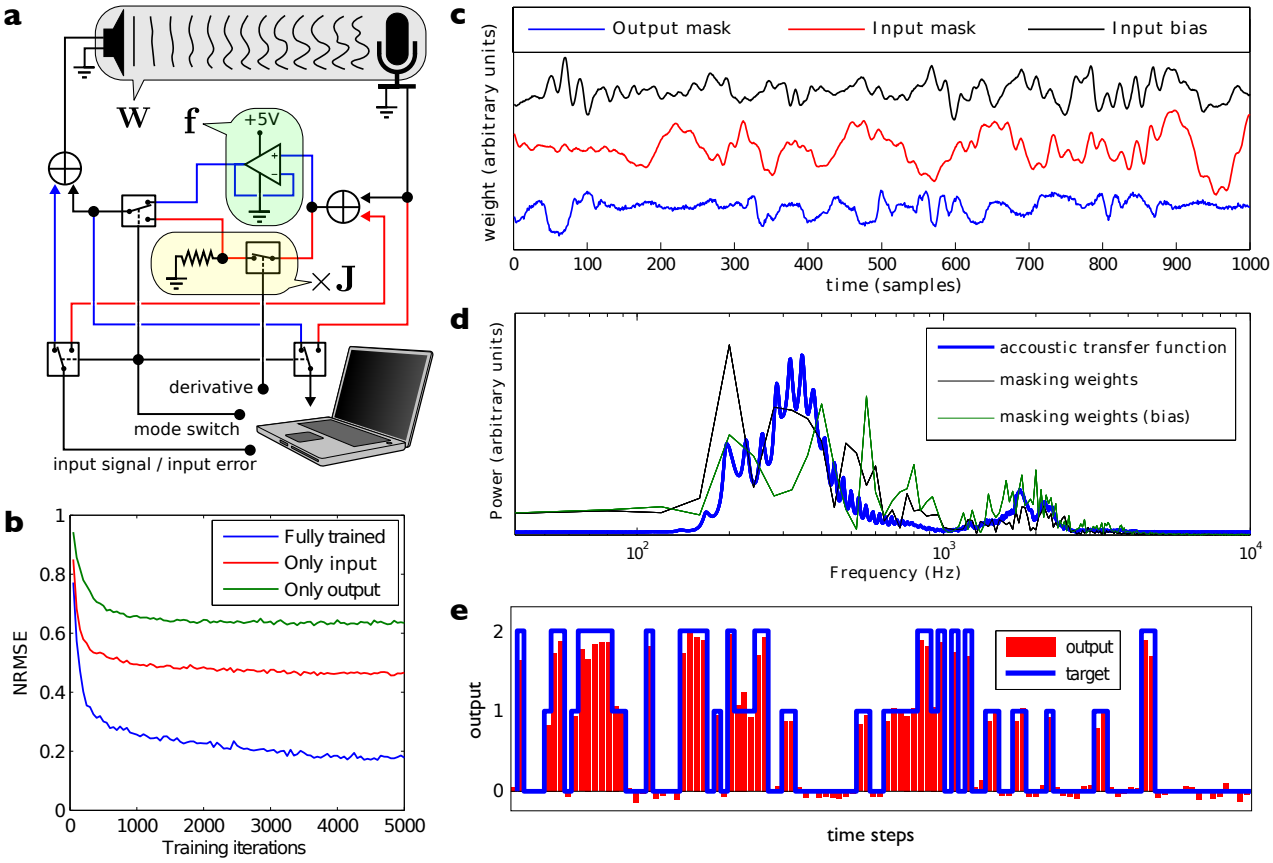


**Figure 2:** Results of the acoustic BPTT experiment. **a:** Schematic depiction of the electric circuit of the experiment. The grey box at the top represents the speaker, tube and microphone. The blue circuit lines are only in use during the forward propagation, and the red lines are in use during the backpropagation phase. The forward or backward mode can be toggled by a logic signal that comes from the PC, and which controls three analog switches in the circuit. In the forward mode, the nonlinear feedback is implemented by an op-amp voltage follower which cuts off the signal at 0V (green box), implementing a linear rectifier function. In the backprop mode, the multiplication with the Jacobian is implemented by a fast analog switch that either outputs zero, or transmits the signal. **b:** NRMSE as a function of the number of training iterations for three cases. That where both the output and input masks are trained, and those where only one of each is trained. **c:** Resulting input and output masks after training. **d:** Power spectrums of the input masks and the power spectrum of the transmission of the speaker-tube-microphone system. **e:** Example of the network output vs. the target.

cated at 0V (such that only positive voltages can pass), implementing a so-called linear rectifier function, which will function as the system non-linearity $\mathbf{f}$. The linear rectifier function is a currently popular activation function in neural architectures [7]. Feedback is implemented by adding this signal to the external input signal. One strong advantage of the linear rectifier function is that its derivative is a binary signal: equal to one when the signal is transmitted, and equal to zero when it is cut off. This means that multiplication with the Jacobian is equivalent to either transmitting the feedback signal unchanged, or setting it to zero, which can be easily implemented using an analog switch. For a more detailed explanation of the relation between the acoustic system and the general case described by Equation 5 we refer to the methods section.

We will use the physical BPTT setup to train input and output masks. Note that the output masks are trained simultaneously with gradient descent, but they do not require the physical backpropagation. Both input encoding and output decoding happen on an external PC.

We tested the setup on an academic task which combines the need for nonlinearity and memory. The input time series $q_i$ is scalar and consists of a series of i.i.d. integers from the set $\{0, 1, 2\}$, which are encoded to an acoustic signal as described above. The desired output time series $y_i$ is defined as

$$y_i = q(i - q(i)),$$

i.e., the task consists of retrieving the input with a delay that depends on the current input. The fact that the delay is variable makes that the task is nonlinear. i.e., it cannot be solved by any linear filtering operation on the input.

For details concerning the experiments we refer to the supplementary material. A schematic depiction of the setup and the main results of the experiments are shown in Figure 2. In Figure 2e we show a comparison between the system output and the target, indicating that the system has learned to solve the task successfully. We show the evolution of the normalised root mean square error (NRMSE) during the training process in Figure 2b. To make sure that the physi-

cal backpropagation works as intended we have ran two additional tests where we either only train the output masks (the classical RC setup) or only the input masks, keeping the other random and fixed. As can be seen in Figure 2b, training only output or input masks in both cases reduces the NRMSE. Note that if the input masks are being trained while the output masks are kept fixed and random, all adaptations to the system parameters are exclusively from the error signal that has been propagated through the system in the form of sound. On top of this, the input mask training needs to find a solution that works with a completely random instantiation of the output mask. The fact that it can achieve this at least to some degree (NRMSE $\approx 0.47$) demonstrates that physical BPTT works as intended.
The input and output masks are shown in Figure 2c. In Figure 2d we show the power spectra of the input masks (and consequently the power spectra of the signals sent into the system) compared to the power spectrum of the system transmission. Clearly, the acoustic parts of the full system (speaker-tube-microphone) only transmit certain frequency bands (a.o. the resonance frequencies of the tube are visible as a set of peaks.), and the input masks seem to have learned to match this spectrum. Note that at no point in training or testing we ever required a model of the acoustic part of the system. The impulse response matrices of the system do not need to be known for the backpropagation to work. Also important to mention is that no process on the external computer or in the electric circuitry provides the required memory for the system to store and process the time series. Instead, this happens due to the memory inherent to the acoustic system and to the nonlinear feedback, which the training process learns to utilise. This means that at any point in time, information about past inputs exists solely as acoustic waves traveling through the tube.

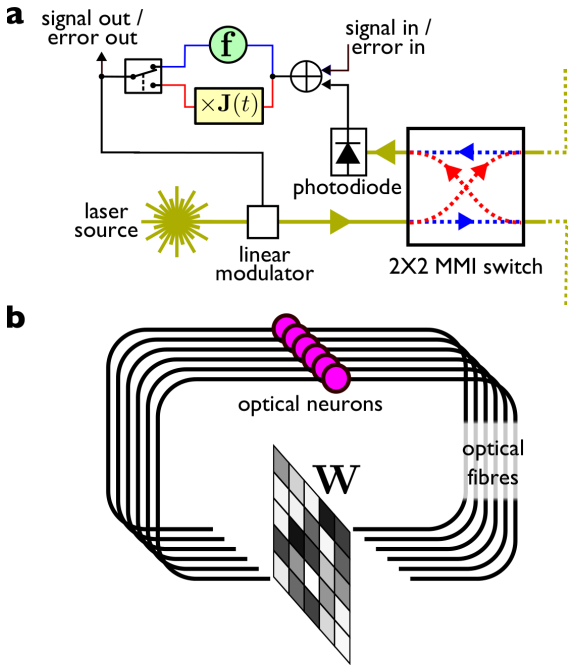


**Figure 3: a:** Schematic depiction of an optical neuron, where pathways that are exclusive in the forward or backward mode are depicted in blue or red, respectively. It is largely similar to the electronic circuit in the acoustic setup, with the difference that the input now enters the system before the nonlinearity. The 2×2 MMI switch allows light to travel either forwards or backwards through the fibre network (fibres depicted by yellow lines). **b:** Depiction of a network of optical neurons, Each purple circle represents a neuron. They send their output signals through optical fibres (which also incorporate delay) to an optical matrix-vector multiplier which multiplies with a matrix $\mathbf{W}$ in the forward direction and $\mathbf{W}^\mathsf{T}$ in the backwards direction. Note that each neuron also has incoming and outgoing connections to external hardware to send input and record output (not depicted).

## A conceptual optics implementation

The described acoustic implementation can be extended to a larger and faster system relatively easy. For instance, one could use ultrasound for higher data transmission (combined with a medium with a high speed of sound). One could use piezoelectric transducers that can both emit and receive signals. Sound propagation is in that sense an interesting platform for a physical analog neural network. The most attractive medium, however, but also more technologically challenging, would be light. Light can transport information at a very high speed, and unlike sound waves, it can be easily guided through fiber optics and integrated photonic waveguides. Just like sound waves, light transmission is reciprocal, making it possible to perform error backpropagation physically on the system. Indeed, using light as a medium for neuro-inspired processing has been studied extensively in the past [4, 27]. These examples primarily exploit parallel processing that happens when light travels through a (often holographic) medium. In our case, we would like to exploit not only parallelism, but also the time delays that are inherent to traveling light.
As a proof of concept we will propose a circuit to perform physical backpropagation optically, partially inspired by the ones described in [16, 20]. Delays are physically implemented by means of long optical fibres. For this example we wish not just to train the input masks, but we wish to control the way in which the signals are mixed as well. Concretely, if we have an $N$-dimensional state $\mathbf{a}(t)$, we wish to optically implement a mixing matrix $\mathbf{W}$ of size $N \times N$, such that the mixing matrix $\mathbf{W_{aa}}(t) = \delta(t-D)\mathbf{W}$, where $D$ is the delay introduced by the optical fibres. Setups for computing matrix-vector products optically have been experimentally demonstrated in the past [8, 26], and here we will assume it is possible to perform them all-optically (see methods). Note that we do not need very high accuracy for neural network applications, as detection of the signal will already be noisy.
Similar to the acoustic example, we conceive an electro-optical 'neuron', which sends and receives optical signals, applies a non-linear function and multiplies with the Jacobian. We encode the state $\mathbf{a}(t)$ as light intensity. Each neuron will have a fixed-power laser source, which can be modulated between a minimal and maximal value. The nonlinear function can be simply implemented by truncating the signal between the minimum and maximum intensity levels of the

laser, which conveniently makes full use of the signal range. The Jacobian of such a function is a again a simple binary function, which can be readily implemented electronically. Finally, we can send in the light into the optical circuit either in the forward and backwards direction by using a $2 \times 2$ MZI optical switch

The final setup we simulate 20 optical nodes, connected with fibres that incorporate a delay of 109 measurement samples. We add noise to the measured intensity, which gives a signal-to-noise ratio of about 18 dB. For the backpropagation phase, we encounter the problem that the truncation at the minimal and maximal intensity still occurs. If we use a small amplitude for the input error signal, such that the truncation levels are not reached, the signal may be overpowered by the system noise. On the other hand, if we make the magnitude too large, the error signal would be incorrectly truncated. We used a scaling factor that balances these two effects. Some truncation still occurs, but it did not significantly impair the training process.

We applied the model of this system on a realistic phoneme recognition task, which is part of speech processing and hence a typical example of a problem that can be solved using recurrent neural networks. We used the often-used TIMIT dataset, which is a speech corpus where the phonemes are directly labeled on the speech signal itself. As an error measure, we use the frame error rate, which is the fraction of frames (time steps) in the input signal that have been mislabeled. More details can be found in the methods section. We found that the system is capable of achieving an error rate of 29.6%. Comparing this to other results in literature we find that we don't perform particularly well, but are in the same ballpark as some other established machine learning techniques (see for instance an overview of results on frame error rate in [14], where the results range from 25% to 39%). When we compare to the RC paradigm, similar results are only obtained using extremely large reservoirs, with up to 20,000 nodes [23].

The numerical simulation took roughly 48 hours to finish on a high end processor, and about 95% of this time was spent on simulating the network, where the rest is mainly spent on computing the parameter gradients. We can roughly estimate how fast the physical setup can be. If we assume that input and output is performed at a speed of one Gigasample per second, which is achievable in commercially available DACs and ADCs, the total time required for gathering all the physically measured data for training the system would take only about 100 seconds, a speedup of three orders of magnitude. Notice that this means that in the current way the training was performed, almost all computation time would be spent on an external processor that computes parameter gradients from the recorded batches of data. Gradient calculations only involve matrix-matrix multiplications, which can again be greatly accelerated using massively parallel processors such as GPUs. Also, the speed at which we can obtain data would make second order methods, such as the Hessian-free approach [19] highly attractive. This method requires only a small amount of updates, but it does require that gradients are computed on as much data as possible. Finally, the speed of data acquisition can help us tackle noise induced problems. If there is a substantial amount of noise present in the measurements, we can simply repeat the measurement several times and average out the results.

# Discussion

In this paper we have proposed a framework for using reciprocal physical dynamic systems with nonlinear feedback as analog recurrent neural networks. We demonstrated that the error backpropagation algorithm, which efficiently optimises a recurrent neural network, can be implemented physically on the same system, thus greatly reducing the necessary computations required for the optimisation process, hence speeding up the training process manifold.

We have experimentally verified the proposed system using an acoustic setup, where signal processing is performed using sound waves travelling through a plastic tube, and a single nonlinear electronic element provides feedback. Importantly, detailed knowledge of the linear part of the dynamic system is not required in order for the backpropagation to work. No matter how complex the transformation of this part of the system is, the reciprocity makes sure that the error backpropagation happens correctly.

To our knowledge, the acoustic example shown in this paper is the first demonstrated non-digital example of a neural architecture where both the processing itself as the error backpropagation happens in a completely dynamic fashion. This research therefore provides a big step towards a completely novel form of processing which relies far less on digital systems, and can directly utilise high-dimensional dynamic systems for processing.

The second example we considered is a proof-of-concept design of an electro-optical system. Unlike sound, light can be manipulated and channeled through fibers. We have demonstrated that the error backpropagation algorithm can be used to optimise the linear impulse response matrices of such system, providing a promising framework for the design of fully trainable electro-optical recurrent neural networks.

In the examples we considered in this paper we mainly optimised input encodings. The scalability of this approach is limited, as the encoding still is performed on an external processor. Future research directions therefore should be focused more on the physical implementation itself. Ideally, we would pass the input signal (light or sound) through a programmable medium, i.e., a medium of which we can control the impulse response matrices in high detail. This would offload the computational complexity of the system to the structure of the physical medium that is used, where the previously mentioned bottleneck no longer would occur.

# Methods

## Acoustic experiments

For the acoustic setup we used a data acquisition card, which samples the signal at 40 kHz (well above the maximum frequency that is still

passed through speaker-tube-microphone system). Input and output masks consist of 1000 samples, which means that time series are processed at 40 time steps per second.
We train the input and output masks over the course of 5000 iterations. Initial values for the input and output masks are picked by i.i.d. sampling from a normal distribution with zero mean and variance equal to 0.2 and 0.1, respectively. Each training batch consists of a newly generated time series of 100 instances. This means that each training iteration (forward and backward pass) takes approximately 5 seconds, and the complete training takes about 7 hours.
As absolute scaling of the error signal does not matter for the backpropagation phase, we always normalised and rescaled the error signal before we used it as input. This was to ensure that the signal always remains well above noise levels. Note that this step causes us to lose the absolute magnitude of the gradients. For parameter updates, we therefore normalised the obtained gradients before using them for parameter updates. The learning rate we chose to be 0.25 at the start of the experiment, and we let it linearly decay to zero to ensure convergence at the end.
Mathematically, the full system can be described as follows: Let $s(t)$ be the input signal after the encoding,and $a(t)$ the (scalar) state of the system at time $t$, then

$$a(t) = f\left(\left[W * (s + a)\right](t)\right),$$

where $f(x) = \max(0, x)$, the linear rectifier function, and $W(t)$ is the scalar impulse response of the speaker-tube-microphone system, i.e., it is the signal that would be received by the microphone for a Dirac delta voltage impulse for the speaker. The output of the system is also the state $a(t)$. This means that we can relate the acoustic system to the general case of equation 5 if $\mathbf{W_{sa}}(t) = \mathbf{W_{aa}}(t) = W(t)$, $\mathbf{W_{ao}}(t) = \delta(t)$ (the Dirac delta function), and $\mathbf{W_{so}}(t) = 0$.

## Conceptual photonic setup

The system is described by the following equation:

$$\mathbf{a}(t) = f(\mathbf{W}\mathbf{a}(t - D) + \mathbf{s}(t))),$$

where $\mathbf{W}$ is the matrix, which is implemented optically (see below) The function $f(x)$ truncates the signal between minimal and maximal intensity (which we define as -1 and 1, respectively):

$$f(x) = \begin{cases} -1 & \text{if } x \leq -1 \\ x & \text{if } -1 < x < 1 \\ 1 & \text{if } x \geq 1. \end{cases}$$

For what follows we assume that the coherence length of the used light is very short, such that light intensities add up linearly. In literature, several ways to implement optical matrix-vector multiplication have been discussed. One possible method would be to encode the vector $\mathbf{a}$ as light intensities. We can let the light pass through a spatial array of tuneable intensity modulators and focus the light on the other side to perform a matrix-vector product. This would have the important limitation that all elements of $\mathbf{W}$ would be positive, as light intensities can only be added up. This is a strong drawback, as this would mean that the system can only provide weighted averages of the individual states and cannot use and enhance differences between them. Therefore we envision a different approach where each element of $\mathbf{a}$ is encoded by two light signals with intensities: $\mathbf{k}+\mathbf{a}$ and $\mathbf{k}-\mathbf{a}$, with $\mathbf{k}$ a vector with all elements equal to one. This ensures that the intensities fall into a positive range between 0 and 2 which can correspond to the minimum and maximum output intensity level of each neuron. Now, these constituents are sent to two separate arrays of intensity modulators $\mathbf{W}_1 = \mathbf{K} + \mathbf{W}/2$ and $\mathbf{W}_2 = \mathbf{K} - \mathbf{W}/2$, where $\mathbf{K}$ has all elements equal to one. As all elements of $\mathbf{W}_1$ and $\mathbf{W}_2$ need to be positive, this means that the range of the elements of $\mathbf{W}$ fall within the range $-2$ to 2. If we combine the light signals after the intensity modulation and hence add up the intensities we get:

$$(\mathbf{W}_1 + \mathbf{W}_2)\mathbf{k} + (\mathbf{W}_1 - \mathbf{W}_2)\mathbf{a} = \mathbf{K}\mathbf{k} + \mathbf{W}\mathbf{a}.$$

The first term will simply introduce a constant bias value which we can remove electronically after measuring. The second term now contains the matrix-vector product where the elements of $\mathbf{W}$ can be both positive and negative.

For the simulations we chose a piecewise constant input signal with a fixed sample period. We also made sure that all delays in the system are integer multiples of this period. This assures that we can treat the system as a discrete-time system where each time step corresponds to the state of the system during a single sample period. Note that we only did this to speed up the simulations (if the delays are non-integer multiples, the simulation would need higher time precision), and this is not a requirement for the actual physical setup. We again used the masking scheme, where each masking period consisted of 100 sample periods. The delay $D$ we chose at 109 sample periods. We used a network of 20 optical neurons.
The TIMIT dataset [6] consists of speech, labeled in time by hand. It has a well-defined training and test set, which makes comparison to results in literature possible. Each time step (frame) needs to be labeled with one out of 39 possible phonemes. The input signal consists of a 39-dimensional time series (coincidentally the same dimensionality of the output signal), which encodes the speech signal using MFCCs. For more details on how the data is encoded, please check, e.g., [23].
We trained for 50,000 iterations, where each iteration we randomly sampled 200 sequences of 50 frames from the training set. For the full training process this means that we recorded $50000 \times 200 \times 50$ data points, each containing 100 samples, and we did this twice (for both the forward mode and the backwards mode), leading to $10^{11}$ measurements. At a measuring rate of 1 Gigasample per second this constitutes a total duration of 100 seconds.

We again normalised the gradients. Parameter updates were performed with a learning rate starting at 1, which linearly dropped to zero over the course of the training.

# References


[1] Lennert Appeltant, Miguel Cornelles Soriano, Guy Van der Sande, Jan Danckaert, Serge Massar, Joni Dambre, Benjamin Schrauwen, Claudio R Mirasso, and Ingo Fischer. Information processing using a single dynamical node as complex system. *Nature communications*, 2:468, 2011.

[Bishop, 2006] Bishop, C. M. (2006). *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer.

[2] Daniel Brunner, Miguel C Soriano, Claudio R Mirasso, and Ingo Fischer. Parallel photonic information processing at gigabyte per second data rates using transient states. *Nature communications*, 4:1364, 2013.

[3] Ken Caluwaerts, Michiel D'Haene, David Verstraeten, and Benjamin Schrauwen. Locomotion without a brain: physical reservoir computing in tensegrity structures. *Artificial life*, 19(1):35–66, 2013.

[4] H John Caulfield, Jason Kinser, and Steven K Rogers. Optical neural networks. *Proceedings of the IEEE*, 77(10):1573–1583, 1989.

[5] C. Fernando and S. Sojakka. Pattern recognition in a bucket. In *Proceedings of the 7th European Conference on Artificial Life*, pages 588–597, 2003.

[6] J.S. Garofolo, National Institute of Standards, Technology (US, Linguistic Data Consortium, Information Science, Technology Office, United States, and Defense Advanced Research Projects Agency. *TIMIT Acoustic-phonetic Continuous Speech Corpus*. Linguistic Data Consortium, 1993.

[7] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier networks. In *Proceedings of the 14th International Conference on Artificial Intelligence and Statistics. JMLR W&CP Volume*, volume 15, pages 315–323, 2011.

[8] Joseph W Goodman, AR Dias, and LM Woody. Fully parallel, high-speed incoherent optical method for performing discrete fourier transforms. *Optics Letters*, 2(1):1–3, 1978.

[9] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. Speech recognition with deep recurrent neural networks. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 6645–6649. IEEE, 2013.

[10] H. Hauser, A.J. Ijspeert, R.M. Füchslin, R. Pfeifer, and W. Maass. Towards a theoretical foundation for morphological computation with compliant bodies. *Biological Cybernetics*, 2012. in press.

[11] Michiel Hermans and Benjamin Schrauwen. Training and analysing deep recurrent neural networks. In *Advances in Neural Information Processing Systems*, pages 190–198, 2013.

[12] Michiel Hermans, Benjamin Schrauwen, Peter Bienstman, and Joni Dambre. Automated design of complex dynamic systems. *PloS One*, 9(1):e86696, 2014.

[13] Herbert Jaeger and Harald Haas. Harnessing nonlinearity: predicting chaotic systems and saving energy in wireless telecommunication. *Science*, 308:78–80, April 2 2004.

[14] J. Keshet, D. McAllester, and T. Hazan. Pac-bayesian approach for minimization of phoneme error rate. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 2224–2227, 2011.

[15] Alex Krizhevsky, Ilya Sutskever, and Geoff Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems 25*, pages 1106–1114, 2012.

[16] L Larger, MC Soriano, D Brunner, L Appeltant, JM Gutierrez, L Pesquera, CR Mirasso, and I Fischer. Photonic information processing beyond turing: an optoelectronic implementation of reservoir computing. *Optics express*, 3:20, 2012.

[17] Mantas Lukosevicius and Herbert Jäger. Reservoir computing approaches to recurrent neural network training. *Computer Science Review*, 3(3):127–149, 2009.

[18] W. Maass, T. Natschläger, and H. Markram. Real-time computing without stable states: A new framework for neural computation based on perturbations. *Neural Computation*, 14(11):2531–2560, 2002.

[19] J. Martens. Deep learning via hessian-free optimization. In *Proceedings of the 27th International Conference on Machine Learning*, pages 735–742, 2010.

[20] Yvan Paquot, Francois Duport, Antoneo Smerieri, Joni Dambre, Benjamin Schrauwen, Marc Haelterman, and Serge Massar. Optoelectronic reservoir computing. *Scientific Reports*, 2:1–6, 2012.

[21] I. Sutskever, J. Martens, and G. Hinton. Generating text with recurrent neural networks. In *Proceedings of the 28th International Conference on Machine Learning*, pages 1017–1024, 2011.

[22] Fabian Triefenbach, Kris Demuynck, and Jean-Pierre Martens. Large vocabulary continuous speech recognition with reservoir-based acoustic models. *IEEE SIGNAL PROCESSING LETTERS*, 21(3):311–315, 2014.

[23] Fabian Triefenbach, Azaraksh Jalalvand, Benjamin Schrauwen, and Jean-Pierre Martens. Phoneme recognition with large hierarchical reservoirs. In *Advances in Neural Information Processing Systems 23*, pages 2307–2315, 2010.

[24] Kristof Vandoorne, Wouter Dierckx, Benjamin Schrauwen, David Verstraeten, Roel Baets, Peter Bienstman, and Jan van Campenhout. Toward optical signal processing using photonic reservoir computing. *Optics Express*, 16(15):11182–11192, 2008.

[25] Kristof Vandoorne, Pauline Mechet, Thomas Van Vaerenbergh, Martin Fiers, Geert Morthier, David Verstraeten, Benjamin Schrauwen, Joni Dambre, and Peter Bienstman. Experimental demonstration of reservoir computing on a silicon photonics chip. *Nature communications*, 5, 2014.

[26] Pochi Yeh and Arthur ET Chiou. Optical matrix–vector multiplication through four-wave mixing in photorefractive media. *Optics letters*, 12(2):138–140, 1987.

[27] Francis TS Yu and Suganda Jutamulia. *Optical signal processing, computing, and neural networks*. John Wiley & Sons, Inc., 1992.

# Supplementary material

## 1 Structure

This document is the supplementary material for the paper: *Scalable, Trainable & Dynamic Computing: Error Backpropagation Through Physical Media*. This document consists of three parts. First we motivate the use of the dynamic system we use, secondly we derive a way to define gradients for it, and thirdly we explain how to use time-multiplexing to enhance a low-dimensional system into a high-dimensional one.

## 2 Motivation for the convolutional structure

We will consider systems that are described by the following set of equations:

$$\begin{aligned} \mathbf{a}(t) &= \mathbf{f}\left([\mathbf{W_{sa}} * \mathbf{s}](t) + [\mathbf{W_{aa}} * \mathbf{a}](t)\right) \\ \mathbf{o}(t) &= [\mathbf{W_{so}} * \mathbf{s}](t) + [\mathbf{W_{ao}} * \mathbf{a}](t) \end{aligned} \tag{5}$$

Here, $\mathbf{s}(t)$ is the input signal, $\mathbf{a}(t)$ is the internal state of the system and $\mathbf{o}(t)$ is the output of the system, and we've used the following notation for the convolution operation:

$$[\mathbf{W} * \mathbf{x}](t) = \int_0^\infty dt'\ \mathbf{W}(t')\mathbf{x}(t-t').$$

By specific choices of the input signal $\mathbf{s}(t)$ and the impulse matrices $\mathbf{W_{xy}}(t)$, a wide range of common forms of neural networks are described by this system. First of all, let's consider the case where $\mathbf{W_{aa}}(t) = \mathbf{0}$, $\mathbf{W_{so}}(t) = \mathbf{0}$, and $\mathbf{W_{sa}}(t) = \delta(t)\mathbf{W_s}$, $\mathbf{W_{ao}}(t) = \delta(t)\mathbf{W_a}$. The equation reduces to:

$$\mathbf{o}(t) = \mathbf{W_a}\mathbf{f}\left(\mathbf{W}_s\mathbf{s}(t)\right),$$

which is nothing more than the description of a neural network with one hidden layer. A deep neural network (a multilayered perceptron) with $L$ layers can be constructed also. We choose $\mathbf{W_{so}}(t) = \mathbf{0}$, $\mathbf{W_{sa}}(t) = \delta(t)\mathbf{W_s}$, $\mathbf{W_{aa}}(t) = \delta(t)\mathbf{W_a}$, and $\mathbf{W_{ao}}(t) = \delta(t)\mathbf{W_a}$, with the following respective definitions:

$$\mathbf{W_s} = \begin{bmatrix} \mathbf{W}_0 \\ \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix},$$

$$\mathbf{W_a} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{W}_1 & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{W}_2 & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{W}_3 & \cdots & \mathbf{0} & \mathbf{0} \\ \vdots & \vdots & \vdots & & \vdots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{W}_{L-1} & \mathbf{0} \end{bmatrix},$$

$$\mathbf{W_a} = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} & \mathbf{W}_L \end{bmatrix}$$

If we further assume that $\mathbf{f}$ acts as a scalar function on each of the elements of its argument, we can rewrite the total system as:

$$\mathbf{o}(t) = \mathbf{W}_L\mathbf{f}(\mathbf{W}_{L-1}\mathbf{f}(\mathbf{W}_{L-2}\mathbf{f}(\mathbf{W}_{L-3}\mathbf{f}(\cdots\mathbf{W}_2\mathbf{f}(\mathbf{W}_1\mathbf{f}(\mathbf{W}_0\mathbf{s}(t)))))))),$$

which is exactly the description of a deep neural network. Note that we do not need the transition matrices to be instantaneous (defined by delta functions) per se. If we hold $\mathbf{s}(t)$ fixed, the system may have internal delays, and we would simply need to hold the input for a long enough time for all internal delays to be resolved.

The system can quite simply be transformed into a recurrent neural network. We define the input signal to be piecewise constant with period $\Delta$, and during each period the input is the next instance from a time series $\mathbf{s}_i$. We then simply need to define $\mathbf{W_{so}}(t) = \mathbf{0}$, $\mathbf{W_{sa}}(t) = \delta(t)\mathbf{W_s}$, $\mathbf{W_{aa}}(t) = \delta(t-\Delta)\mathbf{W_a}$, and $\mathbf{W_{ao}}(t) = \delta(t)\mathbf{W_o}$. This system will now act as a recurrent neural network, where each period $\Delta$ the system undergoes a state update.

$$\mathbf{o}(t) = \mathbf{W_o}\mathbf{a}(t),$$

$$\mathbf{a}(t) = \mathbf{f}(\mathbf{W_s}\mathbf{s}(t) + \mathbf{W_a}\mathbf{a}(t-\Delta)).$$

As may be inferred from these three simple examples, far more intricate or unusual neural architectures also fit in the framework described by equation 5. For instance, if we simply keep the input signal constant, i.e., $\mathbf{s}(t) = \mathbf{s}$ and we assume that under this condition, the system will eventually settle into a steady state, the convolution operations simply become matrix multiplications, and the internal state of the system is described by the following implicit equation:

$$\mathbf{a} = \mathbf{f}(\mathbf{W_{sa}}\mathbf{s} + \mathbf{W_{aa}}\mathbf{a}),$$

which almost never occurs in the neural networks literature, yet poses no problem on our physical setup, and such a system can be trained in the same manner as any system described by equation 5.

## 3 Derivation of gradients

There are several ways to derive the gradients using the error backpropagation mechanism. While it is formally possible to start from equation 5 , it is conceptually easier to use a discrete-time approximation of the convolutions, and consider equation 5 to be a highly specific case of a neural network. First we will give a very brief explanation of the backprop algorithm for an arbitrary neural architecture, and next we consider the special case of the convolutional system used in this paper. For an in-depth proof and explanation of the equations we describe below, we refer to, e.g., [Bishop, 2006]

All neural network architectures can be expressed by a directed graph, where each node performs a nonlinear function, and each edge is a weighted connection. When several edges arrive at a node, their values are added up before they enter the node nonlinearity. The error backpropagation algorithm works by first computing the gradient of the output w.r.t, a certain cost function, (which could be called the output error), and next propagating this error backwards through the directed graph by flipping the direction of each edge. Each edge is still a weighted connection, but during the error backpropagation phase, the error at each node is multiplied with the derivative of the nonlinear function (expressing the current sensitivity of the node). Once the error backpropagation phase is completed, each node in the network will be associated with an error value (the gradient of its activation w.r.t. the cost function). The gradient of each weighted connection (edge) now is simply given by the error at the end times the activation at the front. This principle holds true for any kind of directed graph. Instead of considering each node in the graph as a single, scalar function, one can easily expand this to multivariate systems. Now, each node represents a high-dimensional state and optionally a nonlinear function, and each edge represents a linear transformation matrix. Backwards propagation over such an edge is simply a multiplication with the transposed of this linear transformation matrix. Instead of multiplication with the derivative, one multiplies with the jacobian. The gradient w.r.t. the connection matrix is now defined as the outer product of the error vector at the end of the directed edge with the activations at the start.

Suppose we have an input vector $\mathbf{s}$. We can then formally define any neural architecture as follows:

$$\mathbf{x}_i = \mathbf{W}_i^{\mathbf{s}}\mathbf{s} + \sum_k \mathbf{W}_{ik}\mathbf{a}_k$$

$$\mathbf{a}_i = \mathbf{f}_i\left(\mathbf{x}_i\right)$$

$$\mathbf{o} = \sum_k \mathbf{W}_k^{\mathbf{o}}\mathbf{a}_k + \mathbf{W^{os}}\mathbf{s},$$

where $\mathbf{a}_k$ represents the state of the $k$-th node of the neural graph, and $\mathbf{o}$ is the output on which we define a cost function. We assume that a subset of the transition matrices are equal to zero, such that no loops occur in the computational graph. Note that the transition matrices in these equation represent the *incoming* connections of each particular node in the graph. We define an error $\mathbf{e_o}$ as the gradient of the cost function $C$ w.r.t. $\mathbf{o}$:

$$\mathbf{e_o} = \frac{\partial C}{\partial \mathbf{o}}$$

We can now define errors for each node in the graph as follows:

$$\mathbf{e}_i = \mathbf{J}_i^{\mathsf{T}}\left(\sum_k \mathbf{W}_{ki}^{\mathsf{T}}\mathbf{e}_k + \sum_k \mathbf{W}_k^{\mathbf{o}\mathsf{T}}\mathbf{e_o}\right),$$

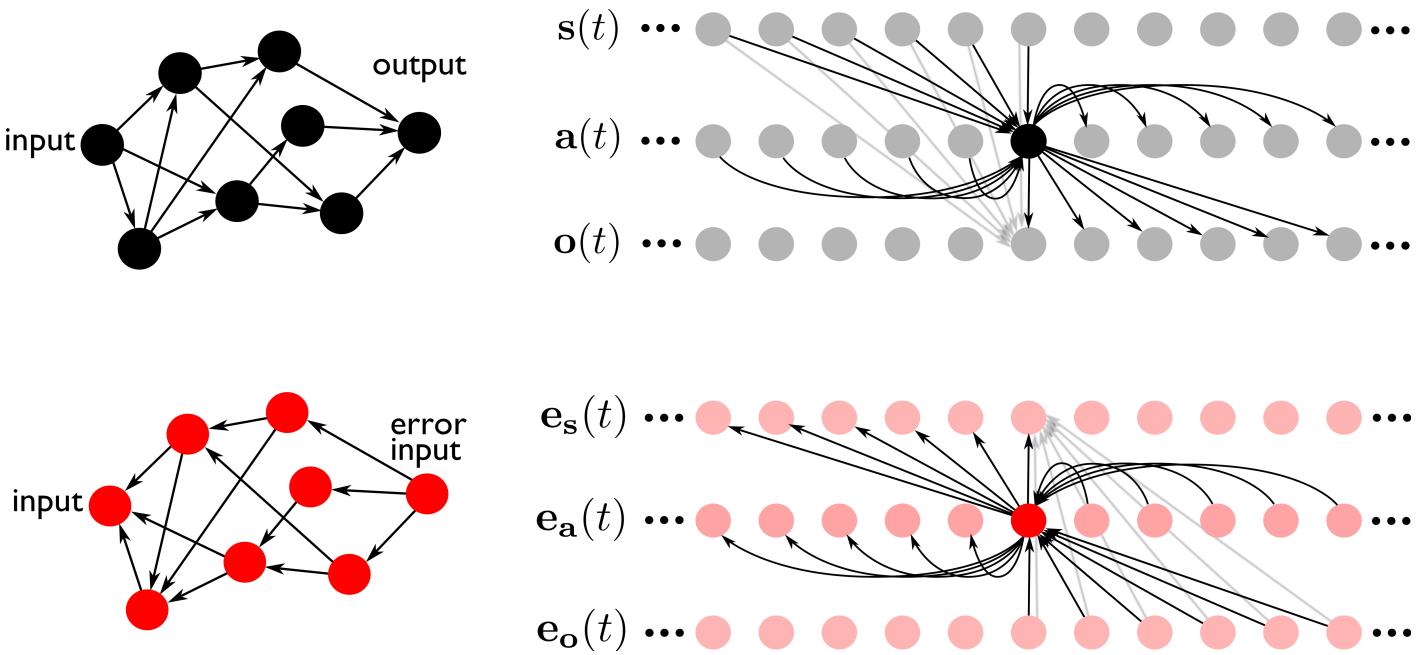


**Figure 4:** Illustration of the graph representation of a neural network architecture. On the left we have depicted the forward and backward graph. Each black circle represents a summation and (optionally) a nonlinear function (except for instance at the output and input nodes). Each arrow represents a linear transformation matrix. During the error backpropagation phase, the direction of all the edges is reversed, and the error is propagated backwards through the network. Once each node is associated with an error, the transition matrices can be updated. On the right we show the graph associated for a time-discretised approximation of Equation 5, where we've shown the incoming and outgoing connections of a single node (and in light grey the connections that directly go from the input to the output, represented by $\mathbf{W_{so}}$). Time goes from left to right. Note that, when the directions of the connections are turned around we again end up with a discrete-time convolution.

with

$$\mathbf{J}_i = \frac{\partial \mathbf{f}_i(\mathbf{x}_i)}{\partial \mathbf{x}_i}.$$

Here, all the transition matrices in these equations represent the *outgoing* connections of each node in the graph, which represents the fact the error propagates *backwards* through the graph. Finally, gradients w.r.t. each transition matrix in the network can be defined as follows:

$$\frac{dC}{d\mathbf{W}_{ij}} = \mathbf{e}_j \mathbf{a}_i^\mathsf{T}, \quad \frac{dC}{d\mathbf{W}_i^\mathbf{s}} = \mathbf{e}_i \mathbf{s}^\mathsf{T}, \quad \frac{dC}{d\mathbf{W}_i^\mathbf{o}} = \bar{\mathbf{e}} \mathbf{a}_i^\mathsf{T}, \quad \frac{dC}{d\mathbf{W}^\mathbf{os}} = \bar{\mathbf{e}} \mathbf{s}^\mathsf{T}$$

Note that these equations hold true for any neural architecture. If some the transition matrices are always equal (which is the case for the convolution we will use), one simply needs to add up the gradients that result from all the single outer products described above.

When making a discrete-time approximation of the convolutions given in equation 5, it can be reduced to a directed graph with the properties described above. First of all, we use the following approximation for the convolution:

$$[\mathbf{W} * \mathbf{y}](t) \approx \delta_t \sum_k \mathbf{W}(k\delta_t)\mathbf{y}(t - k\delta_t),$$

which converges to the proper convolution for $\delta_t \to 0$. We have represented the associated computational graph in the bottom of Figure 4, where we've depicted the incoming and outgoing connections of a single node (time instance) of the state vector $\mathbf{a}(t)$. Note that the connections are shared over time, i.e., the incoming and outgoing connection matrices for different time instance of $\mathbf{a}(t)$ are the same. Also the incoming and outgoing connection matrices that connect the nodes representing $\mathbf{a}(t)$ are symmetric, in the sense that the transition matrix connecting the node representing $a(t-d)$ to $a(t)$ is always equal to the transition matrix connecting node $a(t)$ to $a(t+d)$.

We will use the graph representation of the time-discretised approximation of Equation 5. We start by defining $\mathbf{e_o}(t)$ as the derivative of the cost function w.r.t. $\mathbf{o}(t)$, We time discretise $\mathbf{e_o}(t)$, and write down the time-discretised version of Equation 5:

$$\begin{aligned}
\mathbf{a}(i\delta_t) &= \mathbf{f}\left(\delta_t \sum_k \mathbf{W_{sa}}(k\delta_t)\mathbf{s}((i-k)\delta_t) + \delta_t \sum_k \mathbf{W_{aa}}(k\delta_t)\mathbf{a}((i-k)\delta_t)\right) \\
\mathbf{o}(i\delta_t) &= \delta_t \sum_k \mathbf{W_{so}}(k\delta_t)\mathbf{s}((i-k)\delta_t) + \delta_t \sum_k \mathbf{W_{ao}}(k\delta_t)\mathbf{a}((i-k)\delta_t) \qquad (6)
\end{aligned}$$

We can now simply apply the previously defined general equations for backpropagation through a neural architecture. By switching the outgoing with the incoming connections we can write down:

$$\mathbf{e_a}(i\delta_t) = \mathbf{J}^\mathsf{T}(i\delta_t)\left(\delta_t \sum_k \mathbf{W}_{\mathbf{ao}}^\mathsf{T}(k\delta_t)\mathbf{e_o}((i+k)\delta_t) + \delta_t \sum_k \mathbf{W}_{\mathbf{aa}}^\mathsf{T}(k\delta_t)\mathbf{e_a}((i+k)\delta_t)\right) \qquad (7)$$

When $\delta_t \to 0$, this reduces again to the continuous-time equation for the backpropagation as provided in the main text. When $\mathbf{s}(t)$ is a trainable representation of another kind of signal, we will also need to propagate the error to the input signal $\mathbf{s}(t)$, which is given by:

$$\mathbf{e_s}(i\delta_t) = \delta_t \sum_k \mathbf{W}_{\mathbf{sa}}^\mathsf{T}(k\delta_t)\mathbf{e_a}((i+k)\delta_t) + \delta_t \sum_k \mathbf{W}_{\mathbf{so}}^\mathsf{T}(k\delta_t)\mathbf{e_o}((i+k)\delta_t) \qquad (8)$$

This equation too, reduces to the equation given in the main article when $\delta_t \to 0$. If errors are computed over a time span $t \in \{0 \cdots T\}$, we can define gradients as

$$\begin{aligned}
\frac{dC}{d\mathbf{W_{sa}}(t)} &= \int_0^{T-t} dt' \mathbf{e_a}(t'+t)\mathbf{s}^\mathsf{T}(t') \\
\frac{dC}{d\mathbf{W_{aa}}(t)} &= \int_0^{T-t} dt' \mathbf{e_a}(t'+t)\mathbf{a}^\mathsf{T}(t') \\
\frac{dC}{d\mathbf{W_{so}}(t)} &= \int_0^{T-t} dt' \mathbf{e_o}(t'+t)\mathbf{s}^\mathsf{T}(t') \\
\frac{dC}{d\mathbf{W_{ao}}(t)} &= \int_0^{T-t} dt' \mathbf{e_o}(t'+t)\mathbf{a}^\mathsf{T}(t').
\end{aligned}$$

Here the integral stems from the fact that we need to sum op all contributions to the gradients, each from different time steps. This summation becomes an integral in continuous time

# 4 Input encoding through time multiplexing

If we use a physical dynamic system, the input signal we send in is a continuous-time signal. When we wish to process a discrete time series, which is quite typical for many applications, we need to convert the discrete time signal to a continuous time signal. We will do this using the so-called masking principle, first described in [1]. Here we will use a slightly generalised definition. First of all we define an encoding which transforms the vector $\mathbf{x}_i$, the $i$-th instance of a time series, into a continuous time signal segment $\mathbf{s}_i(t)$:

$$\mathbf{s}_i(t) = \mathbf{s}_b(t) + \mathbf{M}(t)\mathbf{x}_i \quad \text{for} \quad t \in [0 \cdots P]\,,$$

where $P$ is the masking period and $\mathbf{s}_b(t)$ is a bias time trace. The matrix $\mathbf{M}(t)$ are the so-called input masks, defined for a finite time interval of duration $P$. The input signal $\mathbf{s}(t)$ is now simply the time-concatenation of the finite time segments $\mathbf{s}_i(t)$
The output encoding works in a very similar fashion. If there is a time series $\mathbf{y}_i$ which represents the output associated with the $i$-th instance of the input time series, we can define an output mask $\mathbf{U}(t)$. We divide the time trace of the system output $\mathbf{o}(t)$ into segments $\mathbf{o}_i(t)$ of duration $P$. The $i$-th network output instance is then defined as

$$\mathbf{y}_i = \mathbf{y}_b + \int_0^P dt\ \mathbf{U}(t)\mathbf{o}_i(t),$$

with $\mathbf{y}_b$ a bias vector. The process described here is essentially a form of time multiplexing. The backpropagation phase happens very similarly. Suppose we have a time series with instances $\mathbf{e}_i$, which are the gradient of a cost function w.r.t. the output instances $\mathbf{y}_i$. Completely equivalent to the input masking we can now define the error signal $\mathbf{e}_\mathbf{o}(t)$ as a time concatenation of finite time segments $\mathbf{e}^i_\mathbf{o}(t)$:

$$\mathbf{e}^i_\mathbf{o}(t) = \mathbf{U}^\mathsf{T}(t)\mathbf{e}_i.$$

Using this signal as input to the system during backpropagation will provide us with an error $\mathbf{e}_\mathbf{s}(t)$. Exactly equivalent to the output signal, we will now partition $\mathbf{e}_\mathbf{s}(t)$ as a time-concatenation of finite segments $\mathbf{e}^i_\mathbf{s}(t)$. Finally, the gradient w.r.t. the input mask $\mathbf{M}(t)$ is given by:

$$\frac{dC}{d\mathbf{M}(t)} = \sum_i \mathbf{e}^i_\mathbf{s}(t)\mathbf{x}_i^\mathsf{T}.$$

For clarity we have illustrated these principles in Figure 5. Note that in reality, the signals will be sampled using measurement equipment, and we will define $\mathbf{M}(t)$ and $\mathbf{U}(t)$ as discrete-time signals.

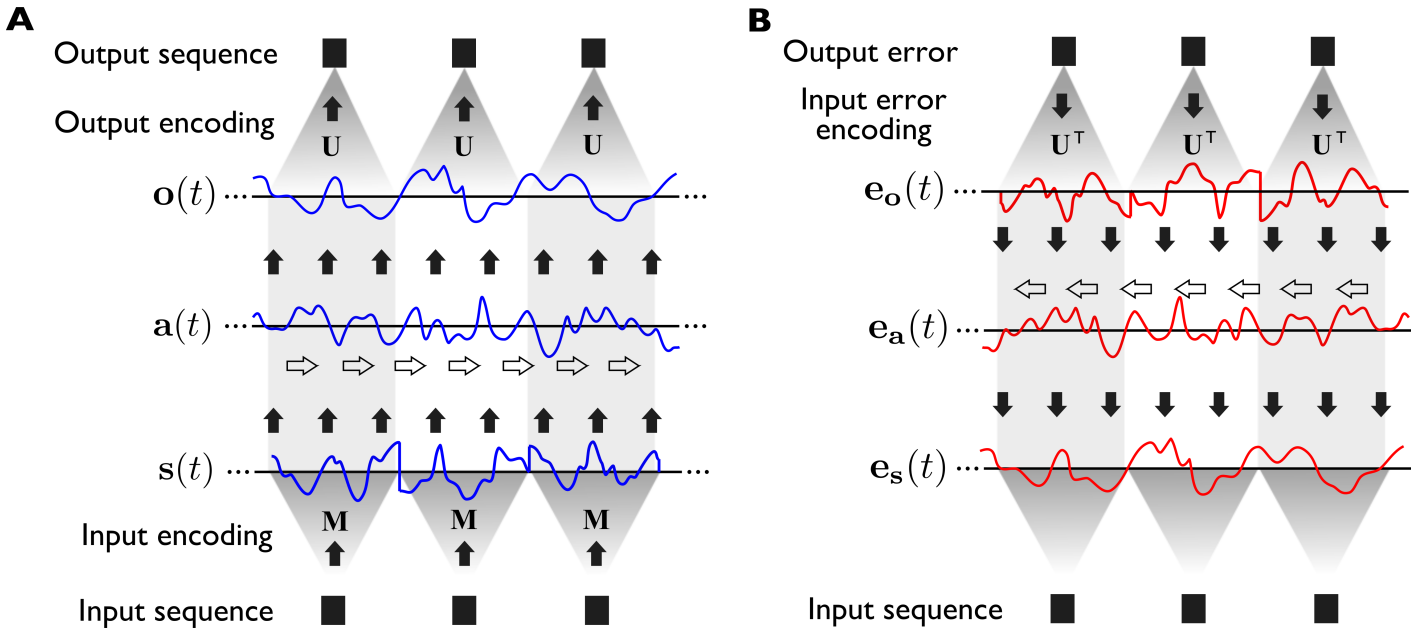


**Figure 5:** Illustration of the masking / time multiplexing principle. **A:** Depiction of the masking principle in the forward direction. At the bottom we see three consecutive instances of an input time series. Each of these is converted into a finite time segment through the masking signals $\mathbf{M}(t)$. These segments are next concatenated in time and serve as the input signal $\mathbf{s}(t)$ for the dynamic system (where time runs according to the white arrows), and which in turn generates an output signal $\mathbf{o}(t)$. The output signal $\mathbf{o}(t)$ is divided into finite length pieces, which are decoded into output instances of an output time series using the output masks $\mathbf{U}(t)$. **B:** The backpropagation process happens in a completely similar manner as in the forward direction. This time, the transpose of the output masks serve as the encoding masks. Finally, the input error signal $\mathbf{e_s}(t)$ is also segmented in time before it is used to determine the gradients w.r.t. $\mathbf{M}(t)$